arXiv:1609.08347v1 [math.ST] 27 Sep 2016

Optimal design of observational studies

Juha Karvanen[1], Jarno Vanhatalo[2,3], Kari Auranen[4,5],
Sangita Kulathinal[2] and Samu Mäntyniemi[6]

[1] Department of Mathematics and Statistics, University of Jyvaskyla, Jyväskylä, Finland, juha.t.karvanen@jyu.fi
[2] Department of Mathematics and Statistics, University of Helsinki, Helsinki, Finland
[3] Department of Biosciences, University of Helsinki, Finland
[4] Department of Health Protection, National Institute for Health and Welfare, Helsinki, Finland
[5] Department of Mathematics and Statistics, University of Turku, Turku, Finland
[6] Department of Environmental Sciences, University of Helsinki, Helsinki, Finland

September 28, 2016

Abstract

Motivated by the universal need for cost-efficient data collection, we propose a unifying framework for optimal design of observational studies. In this framework, the design is defined as a probability measure in the space of observational processes that determine whether the value of a variable is observed for a specific unit at the given time. The framework allows one to optimize decisions on three design dimensions: unit, variable and time. The optimal observational design maximizes the utility defined as a function of model parameters, data and policy decisions. We discuss computational methods that can be used to find optimal or approximately optimal designs. We review a wide variety of design problems that can be cast in the proposed framework, including sample size determination, subsample selection, selection for re-measurements, choice of measurement times, and optimization of spatial measurement networks.

1 Introduction

The cost of data collection is a major factor restricting the accumulation of scientific knowledge. It is therefore of primary interest that the data collection is planned and implemented efficiently. Smart decisions on the study design maximize the precision of parameter estimates with the given budget or, alternatively, minimize costs while ensuring that the goals of the study are met. Despite the importance of the topic, design of observational studies (DOS) is lacking a unified view that would combine methods and insights achieved in different applications fields. We aim to fill this gap by presenting a unifying framework for DOS that allows us to approach a large number of design problems based on the same concepts and principles.

Our framework for DOS encompasses all decisions made on units to be measured, variables to be measured and times of the measurements. In optimal design of observational studies (ODOS), the decisions on some or all of these three design dimensions are optimized by maximizing a predefined utility function under given constraints. The investigator formulates the utility function and the constraints on the basis on the objectives, population, and available resources. In a typical setup, the goal is to maximize the precision of the parameter estimates while the cost of the study is fixed. In the dual problem, the goal is to minimize the cost of the study while the precision of the parameter estimates is fixed.

We define an observational design as a probability measure in the space of observational processes that determine whether the value of a variable is observed for a specific unit at the given time. This definition covers different types of random sampling as well as deterministic designs where the units to be measured are not selected by random sampling but are decided in such a way that the expected utility is maximized. We formulate the principles of ODOS starting from Bayesian decision theory (Raiffa and Schlaifer, 1961; Lindley, 1972; Berger, 1985). ODOS can be viewed as a part of a wider decision theoretic framework where the (monetary) value of information (Raiffa and Schlaifer, 1961; Eidsvik et al., 2015; Lindley, 1985) is used as the leading principle to make decisions about collecting additional data. This extends the scope of applications from scientific research to decision making in business and society.

The proposed framework allows us to consider sample size determination, subsample selection, selection for re-measurements, choice of measurement times and optimization of spatial measurement networks using the same concepts and notation. The benefits of framework are expected to be the largest in complicated situations, such as studies with sequential or multi-stage data collection, because ODOS systematically uses all existing information in the design optimization. If a complex design does not offer benefits over simple random sampling, it is not worth of the effort. In some applications, units are not selected at random but the design is optimized at the unit level. Unit level data are available for design decisions in multi-stage studies and settings where the units are selected from a registry or an existing sample. Unit level optimization of a

design requires computational resources, which may be the major reason why only a few examples of it exist in the literature.

In many applications, survey sampling (Cochran, 1953) can be understood as a synonym for DOS. However, survey sampling has traditionally focused on the estimation of population totals and averages whereas our focus is in statistical modeling estimation and prediction. In survey sampling, statistical modeling is widely applied in small area estimation (Rao and Molina, 2015) but even there the main interest lies in the estimation of the totals by area. Stratified sampling, probability-proportional-to-size sampling and cluster sampling, which aim to improve the precision compared with simple random sampling, can be considered as steps towards ODOS. Our framework can be seen as a generalization of these ideas.

In epidemiology and other fields, specialized designs have been used to boost data collection for rare events such as occurrence of a rare disease. For instance, the widely used case-control design relies on the intuition that if the (disease) cases are rare in the population, sampling cases and controls separately will lead to savings in the sample size compared with simple random sampling. The case-control design (Breslow, 1996) and other intuitively appealing designs (Langholz and Goldstein, 1996; Prentice, 1986) have been successful in practice (Kulathinal et al., 2007; Langholz, 2007) but this does not imply that they would be optimal in a specific problem.

ODOS shares many similarities with optimal experimental design (Kiefer and Wolfowitz, 1959; Pukelsheim, 1993; Atkinson et al., 2007; Goos and Jones, 2011; Chaloner and Verdinelli, 1995). Both are motivated by cost-efficiency and the dual problems for optimization are similar. For the maximization of information, optimal DOE commonly uses criteria such as D-optimality and A-optimality (Atkinson et al., 2007), which are functions of expected Fisher information. The same criteria can be utilized in ODOS. Naturally, experimental and observational designs have the fundamental difference that in the former, the numeric value of a treatment variable for a unit can be decided by the researcher (by means of randomization) whereas in the latter the decision is made whether a unit is measured. Decisions on measurement times and variables to be measured need to be made in both the experimental and observational studies. Some decisions made in experimental design can be actually interpreted as decisions on observational design. For instance, in intervention studies with longitudinal follow-up measurements, the times of the response measurements must be decided and this decision is usually independent on the decision of treatment allocation.

ODOS leads to data missing by design (Wacholder, 1996; Le et al., 1997). As the decision measuring a variable may depend on the observed data (and the prior), the data are missing at random (MAR). Thus, the situation is different from notorious self-selection (Keiding and Louis, 2016) where the inclusion probability may depend on unknown variables. As the data are MAR, the missing data mechanism, that is, the design, may be ignored in direct-likelihood and Bayesian inference (Rubin, 1976). Methods for handling missing data under the MAR assumption can be applied (Little

and Rubin, 2002). The inference is model-based meaning that averages and other descriptive statistics must not be calculated directly from the sample but should be estimated using the model.

The rest of the paper is organized as follows. The notation is introduced in Section 2. The principles of ODOS are formulated in Section 3. Computational methods for finding optimal designs are covered in Section 4. Applications are presented in Section 5. Open problems and future directions are discussed in Section 6.

2 Setting and notation

As a starting point, it is required that the high level objectives of the study are known and the study population, study variables and study period are specified. The study population $\mathcal{I}$ is a set of all observational units $i \in \mathcal{I}$ that can be selected for the study. In many applications especially in medicine and social science, the units are individuals and the study population is discrete and finite $\mathcal{I} = \{1, \ldots, N\}$, where N is the size of the population. In other applications, especially in spatial statistics, the study population (study space) is continuous but has a finite area or volume. The units can then be understood to be locations. The study variables, indexed as $j = 1, \ldots, J$, are variables that can be measured (but are not necessarily measured) during the study. The study period is a union of all time intervals during which measurements can be potentially carried out.

Following the idea of the three design dimensions, let $\{x_{ij}(t)\}$, where i, j and t refer to unit, variable and time, respectively, denote a continuous time process with values in space $\mathcal{X}$. The observational process $r_{ij}(t)$ in space $\mathcal{R}$ is defined as follows

$$r_{ij}(t) = \begin{cases} 1 & \text{if } x_{ij} \text{ is measured at time } t, \\ 0 & \text{otherwise.} \end{cases} \tag{1}$$

The measurements for $x_{ij}(t)$ can be made only at time points in a finite or infinite set $\mathcal{T}^{ij}$. A collection of observational processes $\mathbf{r} = \{r_{ij}(t) : t \in \mathcal{T}^{ij},\ j \in \{1, \ldots, J\},\ i \in \mathcal{I}\}$ is called a measurement plan. The measured data can then be defined using the standard notation for missing data

$$x^*_{ij}(t) = \begin{cases} x_{ij}(t) & \text{if } r_{ij}(t) = 1 \\ \text{NA} & \text{if } r_{ij}(t) = 0. \end{cases} \tag{2}$$

An observational design can be informally defined as a strategy for determining the observational processes $\{r_{ij}(t)\}$ for all i, j and t. More formally, design η is defined as a probability measure on the space of observational processes $\mathcal{R}$. The definition can be compared with the definition of experimental design as a probability measure on

space $\mathcal{X}$ (Chaloner and Verdinelli, 1995). The induced subspace $\mathcal{R}_\eta = \{\mathbf{r} \in \mathcal{R} : \eta(\mathbf{r}) > 0\}$ contains all possible measurement plans in the design η.

The above definition of design allows randomness in the selection of units, variables and measurement times. A design η is deterministic if $|\mathcal{R}_\eta| = 1$, otherwise the design is random. For a deterministic design, set $\mathcal{R}_\eta$ contains only one measurement plan. For a random design, $\mathcal{R}_\eta$ is a set of measurement plans and η defines the probabilities for each of them to be chosen. As an example, consider simple random sampling of units for the measurement of variable j at a fixed time point. In this case, design η is characterized by sample size n_η and selection indicators r_i, $i \in \mathcal{I}$ are sufficient to characterize a measurement plan as the measurement time(s) and the variables to be measured are the same for all units. The set $\mathcal{R}_\eta$ contains all measurement plans $\mathbf{r}$ for which $\sum_i r_i = n_\eta$. The probability of each measurement plan to be chosen is the same $\eta(\mathbf{r}) = 1/|\mathcal{R}_\eta|$. A random design is realized when one measurement plan from set $\mathcal{R}_\eta$ is chosen by random sampling. The null design η_0 is a special case where no new data are collected.

In definition (1), time t refers both to the time in the process $x_{ij}(t)$ and the time when the measurement is actually made. In some cases, these two times may have to be treated separately. We then define

$$r_{ij}(t,s) = \begin{cases} 1 & \text{if the value of } x_{ij}(t) \text{ is measured at time } s, \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

The retrospective collection of medical history is an example of a situation where the two time axes could be used. Time s refers to when the individual was interviewed and the medical records were checked and time t is the time of the recorded event.

We allow the data to have a hierarchical structure with levels $k = 1, \ldots, K$ so that $N(k)$ is the number of clusters at level k and $k = 1$ corresponds to the unit level. The hierarchical structure is specified by an indicator variable z_{ikl} defined as

$$z_{ikl} = \begin{cases} 1 & \text{if unit } i \text{ belongs to cluster } l \text{ on the level } k, \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

The hierarchical structure is assumed to remain constant in time but this assumption can be relaxed if needed.

The general formulation of a design allows us to present a wide variety of problems in the same framework. For a hierarchical data structure, the problem may include deciding the sample sizes for each level of the hierarchy. There is then only one time point and all variables are measured for all units drawn to the sample. In multi-stage studies, the problem may be the selection of units for the measurement of an expensive variable. The measurement times are considered fixed and a hierarchical structure plays no role. In longitudinal studies, the problem may be the choice of the measurement

times. All variables are measured for all units in the study. In replication studies, the problem may be the determination of sample size and selection of variables on the basis of earlier studies. Units are then selected by simple random sampling.

3 Framework for optimal observational design

We assume that our prior knowledge on model parameters θ is expressed by distribution $p(\theta)$ defined in parameter space Θ and our knowledge on the relationship between data and model parameters is described by a model $p(\mathbf{x}|\theta)$. The benefits of the study are evaluated using a utility function and the costs can be measured explicitly in terms of the resources spent.

A general framework for cost-efficient design is Bayesian optimal design (Raiffa and Schlaifer, 1961; Lindley, 1972; Chaloner and Verdinelli, 1995), which maximizes the expected utility obtained from the design conditional on the current information about the problem and the cost structure of the design. The framework is originally presented for experimental design but here we adopt it for observational design. If data $\mathbf{x}_0^*$ have already been observed, the posterior probability distribution $p(\theta|\mathbf{x}_0^*) \propto p(\mathbf{x}_0^*|\theta)p(\theta)$ describes our current knowledge on the model parameters. In a special case where no previous data $\mathbf{x}_0^*$ have been collected, the current knowledge is described by the prior distribution $p(\theta)$ alone. We consider a design η for collecting new data $\mathbf{x}_1^*$ and are interested to know the expected utility achieved by the design. The utility may depend on the model parameters, the new and the current data, the design and the measurement plan realized under the design for collecting the new data. If the data are used for decision making, the utility may also depend on a decision $d \in \mathcal{D}$, where $\mathcal{D}$ is the set of possible decisions. Combining these, the utility function can be written in a general form as $U(d, \theta, \eta, \mathbf{r}, \mathbf{x}_0^*, \mathbf{x}_1^*)$. The expected utility of a design $\mathrm{E}\left(U(\eta|\mathbf{x}_0^*)\right)$ is obtained by conditioning on the existing data $\mathbf{x}_0^*$, marginalizing over measurement plans $\mathbf{r}$ in $\mathcal{R}_\eta$, parameters θ and new data $\mathbf{x}_1^*$, and maximizing over decision d.

The costs of data collection can be incorporated in the utility function. Alternatively, one may define a general cost function $C(\theta, \eta, \mathbf{r}, \mathbf{x}_0^*, \mathbf{x}_1^*)$ that gives costs of the data collection. The costs often depend only on the realized measurement plan, in which case the cost function is simply $C(\mathbf{r})$. The costs are usually defined in terms of money but in some applications it is natural to use time instead of money (Karvanen et al., 2007; Karvanen, 2009).

The new data $\mathbf{x}_1^*$ is a random variable in space $\mathcal{X}_1^*$. Their predictive distribution before the data are collected according to measurement plan $\mathbf{r}$, can be written as

$$p(\mathbf{x}_1^*|\mathbf{r}, \mathbf{x}_0^*) = \int_\Theta p(\mathbf{x}_1^*|\theta, \mathbf{r}, \mathbf{x}_0^*)p(\theta|\mathbf{x}_0^*)\,\mathrm{d}\theta, \tag{5}$$

where $p(\mathbf{x}_1^*|\mathbf{r}, \mathbf{x}_0^*)$ is the model for the observed data. In many applications, $\mathbf{x}_1^*$ and $\mathbf{x}_0^*$ may be assumed to be conditionally independent given θ. However, this does not

hold e.g. for some time series or spatial problems. On the basis of the new data, our knowledge on the model parameters will be updated to a posterior distribution $p(\theta|\mathbf{x}_1^*, \mathbf{x}_0^*, \mathbf{r}) = p(\theta|\mathbf{x}_1^*, \mathbf{x}_0^*)$ where the conditioning on $\mathbf{r}$ is not needed because, according to definition (2), the same information can be deducted from $\mathbf{x}_1^*$.

Next we write the expected utility in an explicit form. First we consider deterministic designs and then generalize to random designs. For deterministic designs marginalization over measurement plans is not needed because the design contains only one measurement plan, i.e., $\mathcal{R}_\eta = \{\mathbf{r}\}$. The optimal deterministic design η^{opt} is then equivalent to choosing the measurement plan that maximizes the expected utility

$$\begin{aligned}&\mathrm{E}\left(U(\eta|\mathbf{x}_0^*)\right) = \mathrm{E}\left(U(\mathbf{r}|\mathbf{x}_0^*)\right) = \\ &\int_{\mathcal{X}_1^*} \left[\max_{d\in\mathcal{D}} \int_\Theta U(d,\theta,\mathbf{r},\mathbf{x}_1^*,\mathbf{x}_0^*) p(\theta|\mathbf{x}_1^*,\mathbf{x}_0^*)\,\mathrm{d}\theta\right] p(\mathbf{x}_1^*|\mathbf{r},\mathbf{x}_0^*)\,\mathrm{d}\mathbf{x}_1^*. \qquad (6)\end{aligned}$$

Constraints on decisions and design can be defined either by incorporating them into the utility function or by restricting the set of possible decisions and designs. The expected costs of a deterministic design are given by

$$\mathrm{E}\left(C(\eta|\mathbf{x}_0^*)\right) = \mathrm{E}\left(C(\mathbf{r}|\mathbf{x}_0^*)\right) = \int_{\mathcal{X}_1^*}\int_\Theta C(\theta,\mathbf{r},\mathbf{x}_1^*,\mathbf{x}_0^*) p(\theta|\mathbf{x}_1^*,\mathbf{x}_0^*)\,\mathrm{d}\theta p(\mathbf{x}_1^*|\mathbf{r},\mathbf{x}_0^*)\,\mathrm{d}\mathbf{x}_1^*. \qquad (7)$$

Ifthe costs depend only on the measurement plan, the integration can be omitted and the costs are fixed already before the data collection. For the null design, it is natural to assume that $C(\eta_0|\mathbf{x}_0^*) = 0$.

If η is a random design, the expected utility can be expressed as a weighted average over all possible measurement plans

$$\mathrm{E}\left(U(\eta|\mathbf{x}_0^*)\right) = \int_{\mathbf{r}\in\mathcal{R}_\eta} \eta(\mathbf{r})\mathrm{E}\left(U(\mathbf{r}|\mathbf{x}_0^*)\right),$$

where $\mathrm{E}\left(U(\mathbf{r}|\mathbf{x}_0^*)\right)$ is defined in equation (6). The expected cost can be written in a similar way.

The frequentist approach to ODOS can be presented similarly to criterion (6) with the exception that, instead of the prior and posterior distributions $p(\theta|\mathbf{x}_0^*)$ and $p(\theta|\mathbf{x}_0^*, \mathbf{x}_1^*)$, only fixed parameter values are considered. If $\hat{\theta}_0$ stands for the initial estimate of θ based on current data $\mathbf{x}_0^*$, the expected utility of a deterministic design can be written as

$$\mathrm{E}\left(U(\eta|\mathbf{x}_0^*)\right) = \mathrm{E}\left(U(\mathbf{r}|\mathbf{x}_0^*)\right) = \int_{\mathcal{X}_1^*} \left[\max_{d\in\mathcal{D}} U(d,\hat{\theta}(\mathbf{x}_0^*,\mathbf{x}_1^*),\mathbf{r},\mathbf{x}_0^*,\mathbf{x}_1^*)\right] p(\mathbf{x}_1^*|\mathbf{r},\hat{\theta}_0)\,\mathrm{d}\mathbf{x}_1^*, \qquad (8)$$

where $\hat{\theta}(\mathbf{x}_1^*, \mathbf{x}_0^*)$ is the point estimate of θ based on all the data available after the new study. Both the Bayesian (6) and frequentist criteria (8) require an integration over the

future data, and maximization over the decision alternatives. Integration with respect to the prior and the posterior distribution are naturally needed only when applying the Bayesian criterion.

Here we restrict to problems where the design costs are deterministic and consider the following three types of optimization:

(A) Optimize the design of a study given the maximum cost C_0

$$\max_{\eta} \mathrm{E}\left(U(\eta|\mathbf{x}_0^*)\right) \text{ given } C(\eta|\mathbf{x}_0^*) \leq C_0,$$

where the utility function is assumed not to depend on the cost of the design.

(B) Minimize the cost of a study given the lower limit U_0 for the expected utility

$$\min_{\eta} C(\eta|\mathbf{x}_0^*) \text{ given } \mathrm{E}\left(U(\eta|\mathbf{x}_0^*)\right) \geq U_0,$$

where the utility function is assumed not to depend on the cost of the design.

(C) Maximize the expected utility of a study without fixed limits for expected utility or cost

$$\max_{\eta} \left(\mathrm{E}\left(U(\eta|\mathbf{x}_0^*)\right)\right),$$

where the utility function is assumed to include also the cost of design.

Problems (A) and (B) are dual problems. In problem (A), the researcher has a fixed budget and the research question is determined. The researcher selects the design that gives maximal information on the research question, for instance, minimizes the variance of the parameters of the interest.

Problem C is typical in business and societal decision making, where both the outcome following the decision and the costs of the design can be measured on monetary terms. In this case, we can formulate problem C as maximization of the (monetary) value of information (Raiffa and Schlaifer, 1961; Lindley, 1985; Eidsvik et al., 2015). To do this, we define utility as a function of the value of decision, $U(v)$, where $v(d, \theta, \eta, \mathbf{r}, \mathbf{x}_0^*, \mathbf{x}_1^*)$ is the monetary value of decision d for each combination of parameter values, data, design and the realized measurement plan. The value of information is then the expected monetary value of observing the outcome from design η, that is, the price V such that (Lindley, 1985; Eidsvik et al., 2015)

$$\int_{\mathcal{X}_1^*} \left[\max_{d \in \mathcal{D}} \int_{\Theta} U\left(v(d, \theta, \mathbf{r}, \mathbf{x}_1^*, \mathbf{x}_0^*) - V\right) p(\theta|\mathbf{x}_1^*, \mathbf{x}_0^*) \, \mathrm{d}\theta \right] p(\mathbf{x}_1^*|\mathbf{r}, \mathbf{x}_0^*) \, \mathrm{d}\mathbf{x}_1^*$$
$$= \max_{d \in \mathcal{D}} \int_{\Theta} U\left(v(d, \theta, \mathbf{x}_0^*)\right) p(\theta|\mathbf{x}_0^*) \, \mathrm{d}\theta.$$

If utility is a linear function of value, $U(v) = a + bv$, the value of information for a design η reduces to (Raiffa and Schlaifer, 1961)

$$\mathrm{E}\left(U(\eta|\mathbf{x}_0^*)\right) - \mathrm{E}\left(U(\eta_0|\mathbf{x}_0^*)\right), \tag{9}$$

which indicates the increase in the expected utility due to the design when the reference level is set by the null design. Hence, the value of information is the maximum amount that a rational decision maker should be willing to pay for having access to new data collected by design η before choosing decision d. For an eligible design, the value of information should be greater than zero. The null design is the optimal design if no other design is eligible.

Utility functions developed in Bayesian experimental design can be applied in ODOS. Lindley (1956) proposed using the expected Shannon information (Shannon, 1948) of the posterior distribution of the model parameters as the utility function. In normal linear regression models, this choice leads to Bayesian D-optimal designs where the utility function depends on the determinant of the expected Fisher information. The same happens in non-linear models if the posterior is approximated by the normal distribution (Chaloner and Verdinelli, 1995). In ODOS, Fisher information can be written as $I_{\mathbf{x}_1^*}(\theta, \mathbf{x}_0^*)$ where the notation $I_{\mathbf{x}_1^*}$ indicates that the expectation is taken over the new data to be collected. This means that the utility is actually a combination of the expected information of $\mathbf{x}_1^*$ and the observed information of $\mathbf{x}_0^*$. The utility function for observational Bayesian D-optimality can be written as

$$U(\theta, \mathbf{x}_0^*, \mathbf{x}_1^*) = \log\det\left(I_{\mathbf{x}_1^*}(\theta) + I^{\mathrm{obs}}(\theta, \mathbf{x}_0^*)\right) \tag{10}$$

where I^{obs} stands for the observed information.

For clarity of presentation, it was assumed that there are no unintended missing data. Adopting the Bayesian paradigm, the incomplete observations can be viewed as forming part of the more general class of unobservables. They are then technically considered jointly with the model parameters θ, and the utility functions involve an integration over both. This obviously makes the calculations more complicated.

Illustrative examples on the theoretical concepts introduced above are presented in Appendix. The first example considers the selection of individuals for which the response variable should be measured in logistic regression. The second example studies optimal sample size allocation in decision making in fishery management.

4 Finding optimal designs

In most cases, optimal observational designs cannot be found analytically because the integrals in criterion (6) cannot be expressed in a closed form. The optimal design may then have to be searched for using numerical methods. However, numerical optimization over all design candidates and measurement plans in criterion (6) is often

computationally intractable. Thus, "optimal designs" are in practice only approximately optimal or the best designs obtained with the used search strategy. Despite this shortcoming, the obtained designs may still be highly cost-efficient as compared with standard (non-optimized) designs.

Ryan et al. (2015a) gives a good overview of the available computational algorithms for Bayesian optimal design. The algorithms are presented for experimental design but they are in general applicable also to observational design because the utility function has the same form. Computational methods are needed for the estimation of posterior distribution, for the estimation or approximation of the utility functions and for searching the design space. Well-known computational methods: Markov chain Monte Carlo (MCMC) (Müller, 1999), sequential Monte Carlo (Drovandi et al., 2014), importance sampling (Cook et al., 2008), Laplace approximation (Ryan et al., 2015b) and approximate Bayesian computation (ABC) (Drovandi and Pettitt, 2013) have been used in the estimation of the posterior distribution and the expected utility.

The search strategies for ODOS slightly differ from those of experimental design because the design space is often formed by a limited number of units. For deterministic designs, we name two basic search strategies as direct search and design search.

The direct search (Karvanen et al., 2009; Reinikainen et al., 2014) is related to the search for exact (discrete) optimal experimental designs and uses heuristic methods, such as the greedy method (Dykstra, 1971) and modified Fedorov method (Cook and Nachtrheim, 1980). In the greedy method applied to problem (A), units are selected one by one as long as the budget allows. The original maximization problem is converted to sequential maximization of designs where only one unit is selected. The approach resembles sequential design but the difference is that the measurements are conducted only after all units are selected, not after each selection as in sequential design.

In the iterative replacement method applied to problem (A), the search starts with an initial design which can be obtained e.g. by the greedy method. The selected units are then considered one by one and replaced by another unit if that increases the value of the utility function. The procedure is iterated until convergence.

The design search (Drovandi et al., 2015) aims to first find the optimal design for the corresponding experimental setup. The ODOS problem is replaced by a DOE problem where the data $\mathbf{x}_0$ represent experimental variables that can be controlled by the researcher. The solution for this problem is an optimal experimental design characterized by measurement points $\mathbf{x}_0^{\text{opt}}$ and their weights. After this, the actual ODOS problem is considered. The observational design is obtained by selecting the units that minimize the distance between the values in $\mathbf{x}_0^*$ and the optimal experimental design $\mathbf{x}_0^{\text{opt}}$. For instance, if the optimal experimental design is a two-point design, half of the observational units are selected from the vicinity of each design point. As each unit can be selected only once, the observational design will differ from the optimal experimental design. The approach is expected to work well if the number of units to be selected is small, the number of units available is large, and the number of variables

is small to avoid the curse of dimensionality in the distance measure. Otherwise it may happen that all candidate units are far from the optimal experimental design.

The computational requirements are usually lighter for random designs than for deterministic designs. Although the equations for random designs are mathematically more complicated than the equations for deterministic designs, it is possible to replace the enumeration of all possible measurement plans in set $\mathcal{R}_\eta$ by random sampling from the set.

5 Applications

In this section, we present some applications of ODOS that are related to epidemiology, fishery management and spatial measurement networks. The applications have been originally introduced using their own notation but we review them using the unified framework and notation defined in Sections 2 and 3.

5.1 Subsample selection

In the basic two-stage setting variables $\mathbf{x}_0^* = \{x_{ij}\}$, $j = 1, \ldots, J_0$, $i = 1, \ldots, n$ have been measured for the complete sample in the first stage. At the second stage, the problem is to select optimally a subsample of size $n_1 < n$ for which variables x_{ij}, $j = J_0 + 1, \ldots, J$ will be measured. Often, the variables at the second stage are expensive to measure and therefore only a subsample can be measured. For instance, a surrogate endpoint may have been measured at the first stage and a small subsample is selected at the second stage to validate the surrogate measurements.

A measurement plan can be characterized by selection indicators $\mathbf{r} = (r_1, \ldots, r_n)$, where $r_i = 1$ if the variables x_{ij}, $j = J_0 + 1, \ldots, J$ are to be measured for unit i. The data collected at the second stage is denoted by $\mathbf{x}_1^* = \{x_{ij}^*\}$, $j = J_0 + 1, \ldots, J$, $i = 1, \ldots, n$, where $x_{ij}^* = \text{NA}$ if $r_i = 0$. If a deterministic design is used, the design problem of type A takes the form

$$\max_{\mathbf{r}} \int_{\mathcal{X}_1^*} \left[\int_\Theta U(\theta, \mathbf{x}_1^*, \mathbf{x}_0^*) p(\theta | \mathbf{x}_1^*, \mathbf{x}_0^*) \, \mathrm{d}\theta \right] p(\mathbf{x}_1^* | \mathbf{r}, \mathbf{x}_0^*) \, \mathrm{d}\mathbf{x}_1^*$$

$$\text{given the constraint } \sum_{i=1}^{n} r_i = n_1. \tag{11}$$

Karvanen et al. (2009) studied the selection of individuals for genotyping in a case-cohort study. The variables available at the time of the selection included phenotypic covariates measured at baseline and a survival outcome measured at the end of the follow-up. Blood samples taken at the baseline had been stored so that genotyping did not require the individual to be alive. A frequentist approach with the D-optimality

criterion was used. In simulation comparisons, the designs were ranked as expected: D-optimal design, case-cohort design and simple random sampling. However, extreme selection, i.e., selecting individuals with the extreme values of phenotypic covariates, worked almost as well as the (approximate) D-optimal design and was recommended as a practical choice. Of note, extreme selection has been a popular design in genetics (Lander and Botstein, 1989; Carey and Williamson, 1991; Darvasi and Soller, 1992; Allison et al., 1998; Van Gestel et al., 2000; Susan and Jack, 2006; Macgregor et al., 2006).

Companies may want to select a subset from their customer database ("big data") for an experiment or a survey. The selection can be formulated as an optimal design problem. Drovandi et al. (2015) proposed the optimal selection of a subset to be analyzed as an alternative for parallel computing in the big data context. As examples, they considered datasets on cancer patients, (simulated) mortgage defaults and accelerometer measurements.

5.2 Selection for re-measurements

In the basic longitudinal setting, measurements are made at pre-specified time points $\mathcal{T}$ for all units and all variables. We consider a variant where only an optimally selected subset of units will be measured. Let $n_1 < n$ be the size of the subset to be measured at time point t_1. The data available before time t_1 can be expressed as $\mathbf{x}_0^* = \{x_{ij}^*(t) : t \in \mathcal{T}, t < t_1\}$, $j = 1, \ldots, J_0$, $i = 1, \ldots, n$. Our notation defines each re-measurement as a new variable and assumes that the indices of the variables are ordered by the measurement time. A measurement plan can be characterized by selection indicators $\mathbf{r}(t_1) = (r_1(t_1), \ldots, r_n(t_1))$, where $r_i(t_1) = 1$ if the variables $j = J_0 + 1, \ldots, J_1$ are to be measured for unit i at time t_1. Writing $\mathbf{r} = \mathbf{r}(t_1)$, the optimization problem (11) is of type A for one time point in a longitudinal study. Before the next re-measurement the same problem is considered again but with updated data.

Reinikainen et al. (2014) and Reinikainen and Karvanen (2016) considered optimal selection of individuals for longitudinal covariate measurements in epidemiological follow-up studies. The interest was to estimate risk factors associated with time-to-event outcomes. Both frequentist (Reinikainen et al., 2014) and approximate Bayesian (Reinikainen and Karvanen, 2016) approaches were applied using D-optimality criterion. Deaths and drop-outs posed special restrictions for the design. The selection preferred older individuals (higher risk of an event) and individuals with extreme covariate values. In the re-analysis of an existing epidemiological data, it was demonstrated that the sampling size with ODOS can be 10–25% smaller than under simple random sampling.

5.3 Sample size determination

Sample size determination is a special case where the design decision concerns the number of units to be measured and the units are selected by random sampling. The problem has been studied by many authors in settings where available resources need to be allocated between either multiple stages or multiple subgroups. Closed-form solutions are often available in the frequentist approach. Model-based survey sampling (Chambers and Clark, 2012) can be seen as a DOS problem where the decisions concern sample sizes.

We mention some examples where sample size determination has been approached as an optimal design problem. Reilly (1996) considered optimal sampling strategies for epidemiological two-stage studies. McNamee (2002) and Wruck et al. (2006) studied optimal designs of two-stage studies for the estimation of the sensitivity and the specificity of a diagnostic test. Bekmetjev et al. (2012) proposed a cost-efficient resampling design for the situation where a fraction of the sample is classified by the same imperfect method twice. Rezagholi and Mathiassen (2010) reviewed cost-efficient designs from the viewpoint of occupational exposure assessment. They considered multi-stage studies and comparisons of measurement methods and concluded that the reviewed studies had used simplified analytical tools and insufficient economic analyses. Sutton et al. (2007) studied sample size selection in a meta-analytic framework and Nikolakopoulou et al. (2016) extended the approach to network meta-analysis.

In the case of clustered data, the problem is to decide the sample sizes $N(k)$ for hierarchy levels $k = 1, \ldots, K$. The utility function can depend, for instance, on the precision of the estimated population mean. There is only one time point and all variables are measured for the units in the sample. A design is characterized by sample sizes $N(1), \ldots, N(K)$ and a measurement plan is characterized by selection indicators r_i, $i = 1, \ldots, N(1)$. If the cost per the kth cluster level is c_k, the total cost of a design will be

$$\sum_{k=1}^{K} c_k N(k) = \sum_{k=1}^{K} \sum_{l=1}^{L_k} c_k \mathbb{I}\left(\left(\sum_{i=1}^{N(1)} r_i z_{ikl} \right) > 0 \right),$$

where indicator variable z_{ikl} is defined in equation (4), $\mathbb{I}$ denotes an indicator function and L_k denotes the number of clusters on hierarchy level k.

Tokola et al. (2014) presented an application to clustered fishing data. The aim was to collect data on the key properties (length, weight, maturity stage and sex) of the fish catch needed for decisions on fishery policy. The hierarchy had three levels: fishing vessels, fishing journeys and individual fish. Realistic marginal costs (in euro) were set for each additional vessel, journey and individual to be included in the data collection. There were natural upper limits for the number of vessels and the number of journeys. The objective was to minimize the total cost of the data collection when the precision targets for the key properties were derived on the basis of the regulations

by European Union (a type B problem). The authors utilized geometric programming (Boyd et al., 2007) and ended up with an interesting solution where all vessels available should be used but only nine fish per journey should be measured.

Tokola et al. (2011) considered an application to the design of a health coaching study. The problem was to select the number of coaches and the number of subjects per coach in such a way that the power of the study is maximized given the total cost.

5.4 Choice of measurement times

Continuous-time Markov models are often used to describe the dynamics of phenomena e.g. in life sciences. Let $X(t) \in \mathcal{X}$ be a Markov process in a discrete state space $\mathcal{X} = \{1, \ldots, G\}$, governed by transition intensities θ and indexed by continuous time t. When the current state of the process is observed only at discrete time points for each study subject, the question emerges about the optimal timing of the measurements. For illustration, consider measuring the status of bacterial colonisation (Mehtälä et al., 2015a,b) or the presence of a parasitic infection (Nagelkerke et al., 1990).

In practice, measurements can only be made at discrete time points $\mathcal{T} = \{t_0 = 0, \ldots, t_M\}$ within a (possibly very large) sampling frame including n individuals. The measurement plan is now characterised by selection indicators $\mathbf{r} = (\mathbf{r}_1, \ldots, \mathbf{r}_n)$, where $\mathbf{r}_i = (r_i(1), \ldots, r_i(M))$ and $r_i(m) = 1$ if individual i is included in the sample and the state of his/her process is measured at time $t_m \in \mathcal{T}$, and 0 otherwise. Note that this formulation views also set $\mathcal{T}$ as part of the sampling frame since measurements may only be made in a subset of $\mathcal{T}$.

In a frequentist framework, the optimal design is found as the solution to following problem

$$\max_{\mathbf{r}} \int U(\hat{\theta}(\mathbf{x}_1^*), \mathbf{x}_1^*) p(\mathbf{x}_1^* | \mathbf{r}, \hat{\theta}_0) d\mathbf{x}_1^*,$$

with the constraint that $\sum_{i=1}^{n} \sum_{m=1}^{M} r_i(m) = C_0$. Here C_0 is the total number of observations that can be made due to practical and/or budget limitations. Often the same number of observations is to be taken from all individuals included in the sample, leading to an additional constraint that for any individual i *included* in the sample $\sum_{m=1}^{M} r_i(m) = n_1$ for some $n_1 > 0$. The optimal solution will then determine both n_1 and the number of individuals sampled as $N = C_0/n_1$.

Mehtälä et al. (2015a) investigated the optimal sampling interval in a two-state Markov process with equidistant time spacings between consecutive discrete-time observations. The optimal choice of N vs. n_1 (cf. above) was shown to depend on the distribution of the initial condition. In addition, applying a two-stage design with a utility function of type equation (10), the optimal split of a follow-up study into two stages was investigated. The problem of optimal sampling times was addressed for a model with more than two states by Mehtälä et al. (2015b) in a Bayesian framework.

Cook et al. (2008) optimized the measurement times for epidemic processes. Varis et al. (1990) optimized the monitoring strategy (monthly, biweekly or weekly monitoring) in a lake management problem. Ryan et al. (2014) and Ryan et al. (2015b) and applied Bayesian optimal design in the planning of the measurement times in memory retention tests and pharmacokinetic studies. Although Ryan et al. talk about experimental design, the problem of deciding the measurement times has observational nature and can be put under the umbrella of ODOS.

5.5 Spatial problems

In spatial problems, the measurements are spatially indexed so that $x_{ij}(t) = x_j(\boldsymbol{\xi}_i, t)$, where $\boldsymbol{\xi}_i$ is the spatial co-ordinate of the i'th spatial observation unit, and the interest typically lies in spatially indexed latent variables denoted here by $\phi(\boldsymbol{\xi})$. For example, in spatial epidemiology the observations are disease incidences while the latent variables could describe the relative risk for a disease (Elliott et al., 2001; Vanhatalo et al., 2010) and in environmental applications observations may represent, for example, sea surface temperature (Vanhatalo et al., 2016) or rainfall (Sansó and Müller, 1999) whose real values are denoted by the latent variables. The spatial domain is typically either divided into discrete set of spatial areas, in which case there is a finite number of latent variables and possible observational units, or treated continuously, in which case the latent variables are realizations of a continuous latent process and there is an infinite (uncountable) number of possible observational units within a limited study region. Spatial epidemiology, where the spatial domain would consist of administrative regions, is a typical example of the former and environmental applications of the latter. In both cases we can extend the model parameters to include the latent variables, $\phi \subset \theta$.

Designing optimal observation networks in spatial domain is a well-covered subject in statistical literature. Typical examples of ODOS consists of optimal design of observational network for, e.g., weather measurements. Hence, we mention only a few examples to illustrate how earlier work fits in our general ODOS framework. One typical example is presented by Sansó and Müller (1999), who consider a setting where the amount of existing observation stations needs to be decreased by optimally choosing a subset from the existing stations. They use a utility function that depends on the accuracy of the rainfall predictions and the cost induced by the number of stations in that network (a type C problem). In their application the spatial domain is discretized since they consider only a fixed number of stations in predetermined possible locations. Diggle and Lophaven (2006) define the utility as the averaged prediction variance over a spatial area A

$$U(\theta, \eta, \mathbf{r}, \mathbf{x}_1^*, \mathbf{x}_0^*) = \int_{\boldsymbol{\xi} \in A} \mathrm{Var}(\phi(\boldsymbol{\xi}) | \mathbf{x}_1^*, \mathbf{x}_0^*) \, \mathrm{d}\boldsymbol{\xi}$$

where $\mathrm{Var}(\phi(\boldsymbol{\xi}) | \mathbf{x}_1^*, \mathbf{x}_0^*)$ is the posterior variance of the latent process. In their example,

the locations of monitoring stations could be set freely within the area A but their number was fixed (a type A problem). Other examples on optimal design of observation networks are provided by, e.g., Müller et al. (2004) and Müller (2007). Eidsvik et al. (2015) provide a detailed introduction to the theory of value of information based ODOS in the context of spatial decision analysis. They also address several applications related to e.g. mining, forestry and oil drilling problems.

5.6 Choice of variables to be studied

The choice of variables as a design problem occurs at least in two scenarios. In the first one, the objective is to obtain precise predictions at the unit level when there are several potential predictors that could be measured in the new study. The cost of measurement differs across predictors. The second scenario is related to accumulation of scientific information. Researchers should focus their efforts so that the expected scientific impact of a new study is maximized. This has direct implications not only to decisions on the sample size but also on the variables to be measured in the new study. In both scenarios, a measurement plan can be characterized by selection indicators $\mathbf{r} = (r_1, \ldots, r_J)$, where $r_j = 1$ if the variable j is measured for units in the sample. It is possible to extend the problem so that variables to be measured are decided separately for each unit. In this case, a measurement plan is characterized by a selection matrix $\{r_{ij}\}$.

For the second scenario, Karvanen and Sillanpää (2016) proposed a formal approach to deciding which covariates should be measured in a new study in a meta-analytic framework. They compared decision criteria based on conditional power, change of the p-value, change in lower confidence limit, Kullback-Leibler divergence, Bayes factors, Bayesian false discovery rate or difference between prior and posterior expectations. As an illustration, they considered covariate prioritization based on the results of an existing meta-analysis of genome-wide association studies and made suggestions on the genes to be studied further.

5.7 Evaluating value of information

The value of information can be used to find the maximum price the investigator should pay for the additional data. Examples of this can be found from environmental and health care management (Mäntyniemi et al., 2009; McDonald and Smith, 1997; Yokota and Thompson, 2004a,b) and also in the context of earth sciences (Eidsvik et al., 2015). Mäntyniemi et al. (2009) analysed the value of information of resolving uncertainty about a biological hypothesis related to the population dynamics of North sea herring. Two alternative hypotheses were considered: the amount of offspring either approaches an asymptotic carrying capacity as the number of parents increases (Beverton-Holt model (Beverton and Holt, 1957)) or peaks at a certain parental population size and

then declines (Ricker model (Ricker, 1954)). As a result, the maximum price that the fishing industry should be willing to pay for a research that would completely remove the uncertainty about this population dynamic hypothesis was found to be 240 million Norwegian crowns. However, this is was only a 1.6% increase in the expected utility, compared with optimal solution under the prevailing uncertainty. The authors also invoke the concept of "price of overconfidence", the expected loss that would result if ignoring the prevailing uncertainty and optimizing the fishing pressure by assuming that the correct model structure is known.

6 Discussion

We have presented a unifying framework for optimal design of observational studies (ODOS). In this framework, the observational design is defined as a probability measure in the space of observational processes that determine whether the value of a specific variable is to be observed for a specific unit at the given time. The optimal observational design is a sampling strategy that maximizes a utility function under given constraints. We gave examples on the computational methods that can be used to find the optimal design and coined the terms direct search and design search. The framework can be used to describe a wide variety of design problems as demonstrated in the examples of Section 5.

ODOS has many similarities with optimal experimental design (DOE). In both fields, the design problem is solved by maximizing a utility function under a parametric statistical model. The same utility functions can be used for both ODOS and DOE and the computational methods for optimization share many common features. From the theoretical point of view, the main difference is that an observational design is a probability measure in the space of observational processes while an experimental design is a probability measure in the space of actual variables. It follows that in ODOS each unit can be selected only once (at a given time point). In DOE, the design space can be continuous and the same value of a design variable can be repeated as many times as needed. It is also possible to consider hybrid studies that contain elements from both experimental and observational studies. The design space is then a product of the space of experimental variables and the space of observational processes.

ODOS can be criticized for its sensitivity to the assumptions about the parametric statistical model. This criticism is not unique to ODOS but applies to DOE and sample size calculations as well. In the Bayesian framework, model uncertainty can be taken into account using hierarchical structures and hyperpriors. The statistical model used for the design optimization does not bind the hands of the analyst if a better model is found in the analysis phase. Naturally, the design will be sub-optimal for the new model but it is still likely to be better than a non-optimized design if the analysis model and the design model share common features.

Sometimes the researcher may have difficulties in formulating the utility function because the planned study has multiple goals. For instance, there may be several regression parameters to be estimated and it is disputable whether D-optimality, A-optimality or some other compound criterion provides the best way to summarize the overall precision of the estimates. Approaches developed in the field of multiple-criteria decision making (Miettinen, 1999) are potentially applicable for finding and visualizing nondominated solutions (Miettinen, 2014) in ODOS.

ODOS offers many interesting questions for further research. The development of efficient computational methods is one of the most important challenges. This includes both inventing new methods and modifying optimization methods developed for optimal experimental design. Unintended missing data provides an additional challenge for the optimization. Complex study designs and missing data mechanisms can be illustrated using causal models with design (Karvanen, 2015). Finding new applications of ODOS, for instance in epidemiology, environmental sciences and marketing research, is an important topic for the future research. Extensive simulations with realistic parameters are needed for increasing understanding about the potential benefits of ODOS in various applications. It would be also interesting to investigate optimal designs for problems where the researcher can both observe and intervene the system.

Many of the applications presented in Section 5 are based on re-analyses of existing data. These analyses suggest that significant cost savings could be achieved if ODOS was applied in the planning of new studies. The real application of ODOS requires an open-minded principal investigator who wishes to abandon the tradition in order to try a new idea. The bottleneck in many research projects might be the need for a skilled statistician both in the planning and analysis of the study.

Compared with traditional observational designs, ODOS leads to improved precision and cost savings. On the other hand, ODOS is more complicated to implement. It is expected that the benefits will exceed the disadvantages at least in large studies and studies with expensive measurements.

7 Appendix: Illustrative examples

7.1 Selecting individuals to be measured in logistic regression

As an illustrative example, we consider a small observational study where the target is to estimate a logistic regression model where the response Y is explained by covariates z_1, z_2 and z_3. The covariate values are available for 50 units and the design problem is to select ten units for which the response will be measured. This is type A problem.

The model has the form $\text{logit}(P(Y_i = 1)) = a+b_1z_{i1}+b_2z_{i2}+b_3z_{i3}$, where a, b_1, b_2 and b_3 are the parameters to be estimated. Covariates z_1 and z_2 are continous and covariate z_3 is binary and independent on z_1 and z_2. The Bayesian D-optimality criterion (10)

is used, which means that we aim to maximize the logarithm of the determinant of the Fisher information for a, b_1, b_2 and b_3. Parameter b_2 has an informative prior $b_2 \sim \mathrm{N}(1, 0.04)$ and the other parameters have weakly informative priors $a \sim \mathrm{N}(0, 4)$, $b_1 \sim \mathrm{N}(1, 4)$ and $b_3 \sim \mathrm{N}(1, 4)$.

In this small scale example, the (approximately) optimal designs can be found using direct search with Markov chain Monte Carlo (MCMC) simulations or importance sampling. The computation proceeds in accordance with equation (6), where $\theta = (a, b_1, b_2, b_3)$, $\mathbf{x}_0^* = (z_1, \ldots, z_{50})$ and $\mathbf{x}_1^*$ contains the values of Y for ten selected units. The best design among 100 random 10-unit designs is selected as the initial design and modified Fedorov method (Cook and Nachtrheim, 1980) is applied to find the final design. The procedure to estimate the expected utility for each candidate design evaluated during the search utilizes importance sampling. The realizations of θ are generated using the prior as the importance distribution. Independently of this, realizations of Y are generated for the selected units using parameters generated from the priors. For each pair of θ and Y, the logarithm of the determinant of observed Fisher information is evaluted. These values weighted by importance sampling weights are averaged to obtain an estimate for the expected utility.

Figure 1 shows the obtained optimal design. The rare covariate values $z_3 = 1$ are overpresented in the design compared to the full sample. This improves the balance and leads to more precise estimation. Estimated expected utility for the design is about the same as the median from 100 random samples of size 18. As the optimal design has only ten units, this means that ODOS reduces the sample size by about 40% compared with simple random sampling.

7.2 Optimal sample size in decision making

The next example considers a half-fictitious problem from the context of aquaculture. A rearing tank contains 10,000 fish. Unknown proportion θ of these fish have a disease that leads to death if not given a proper treatment. Only the whole tank can be treated, which costs 15,000 euro. The treatment will cure all the infected individuals. After the rearing period, each fish alive yields a profit of 5 euro. Fish can be sampled from the tank with no cost, but diagnosing the disease costs 10 euro/individual. The question of interest is to choose the optimal number of fish to be diagnosed before the decision on the treatment will be done. This is a problem of type C.

The utility function of the farmer is

$$U(\theta, d) = \begin{cases} (1-\theta)10000 \times 5 \text{ euro }, & d = \text{No treatment} \\ 10000 \times 5 \text{ euro } - 15000 \text{ euro }, & d = \text{Treatment.} \end{cases} \tag{12}$$

Based on earlier experience and data x_0^* (omitted for clarity in this example), the current knowledge about the prevalence of the disease is $p(\theta) = \mathrm{Beta}(\theta|3, 3)$. Since

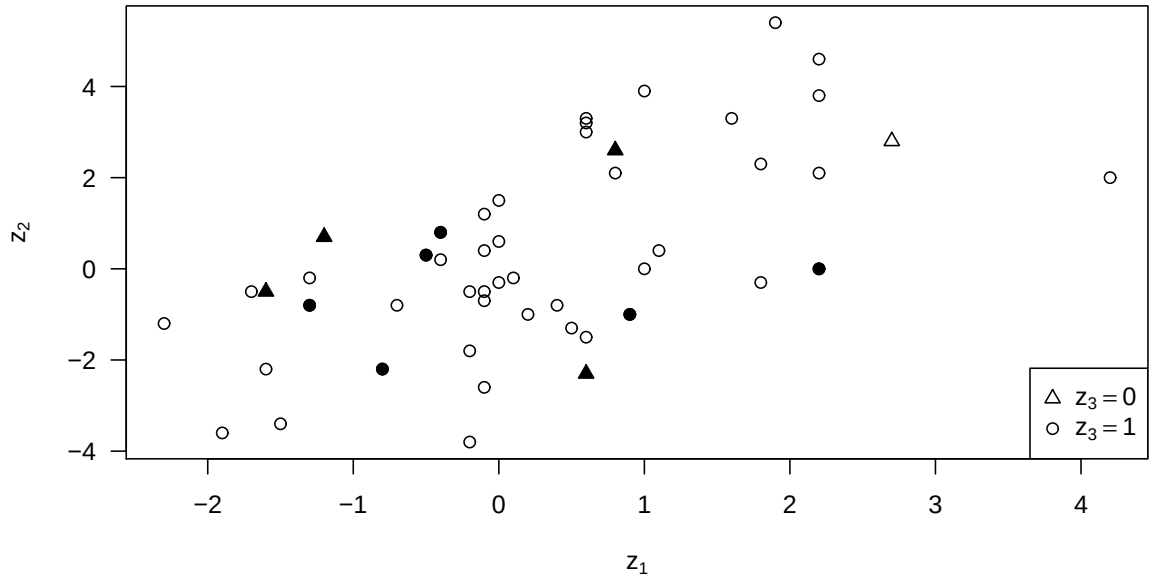

Figure 1: Illustration of the optimal design in the logistic regression example. Ten units (solid circles and triangels) are selected for the study and 40 units (open circles and triangles) are not selected for the study according to the measurement plan.

$\mathrm{E}(\theta) = 3/(3+3) = 0.5$, the optimal decision under the current knowledge is to treat the tank which leads to the utility of 35,000 euro.

Next we evaluate, if the farmer should expect to obtain a higher maximum expected utility after sampling n_η individuals from the tank. For the sake of simplicity, we assume here that fish can be sampled from the tank only one by one with replacement, individuals can not be identified and that infected and clean fish do not aggregate. This leads to a simple binomial model $x_1^* \mid \theta, n_\eta \sim \mathrm{Bin}(x_1^* \mid n_\eta, \theta)$, where x_1^* is the number of infected fish in the sample of size n_η. The expected utility is obtained by first finding the utility for each potential number of infected fish in the sample x_1^* and then averaging over these values according to their respective prior probabilities

$$\mathrm{E}(U(n_\eta)) = \sum_{x_1^*=0}^{n_\eta} \text{Beta-Bin}(x_1^* \mid n_\eta, 3, 3) \max_{d \in \mathcal{D}} \int U(\theta, d) \mathrm{Beta}(\theta \mid x_1^* + 3, n_\eta - x_1^* + 3) d\theta, \qquad (13)$$

where $U(\theta, d)$ is defined in equation (12). Note that $U(\theta, d, n_\eta, x_1^*) = U(\theta, d)$ for all n_η and x_1^*. When calculating $\mathrm{E}(U(n_\eta))$, we integrate over θ and x_1^* and maximize over d.

Figure 2 shows the value of information for all sample sizes from 0 to 100. The value of information is zero for sample sizes 0-4 and positive for larger sample sizes. Thus, the farmer might benefit only from the sample size of five fish and beyond.

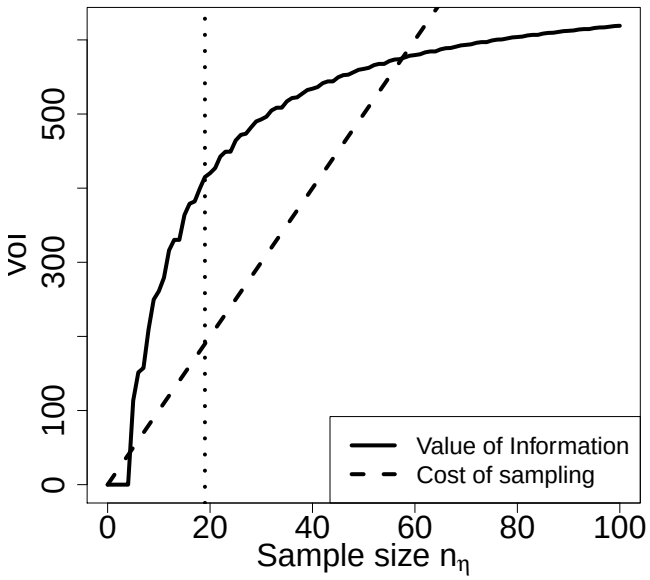

Figure 2: Value of information for the different sample sizes in the fish farmer's problem. The vertical line shows the sample size where the difference between the value of information and the cost of sampling is the highest: this is the optimal sample size.

Now that the farmer knows that sampling at least five is potentially beneficial, she can compare the value of information with the costs of sampling. Since the value of information is the maximum that she should be willing to pay, her task is to calculate the expected utility of each potential sample size by subtracting the cost of the sample size from the value of information. Figure 2 shows that this difference is positive for sample sizes from 5 to 57 and that the difference is the highest for the sample size of 19. The same result is obtained if utility function is directly defined to include the costs.

Acknowledgements

The authors thank Antti Penttinen, Elja Arjas and Mikko Sillanpää for useful comments. Kaisa Miettinen and Markus Hartikainen are acknowledged for discussions on multiple-criteria decision making. JK has been funded by Academy of Finland (grant number 266251). JV has been funded by Academy of Finland (grant 266349) and the Research Funds of the University of Helsinki (decision No. 465/51/2014).

References

Allison, D. B., Heo, M., Schork, N. J., Wong, S.-L., and Elston, R. C. (1998). Extreme selection strategies in gene mapping studies of oligogenic quantitative traits do not always increase power. *Human Heredity*, 48(2):97–107.

Atkinson, A. C., Donev, A. N., and Tobias, R. D. (2007). *Optimum Experimental Designs, with SAS*. Oxford University Press, Oxford.

Bekmetjev, A., VanBruggen, D., McLellan, B., DeWinkle, B., Lunderberg, E., and Tintle, N. (2012). The cost-effectiveness of reclassification sampling for prevalence estimation. *PloS One*, 7(2).

Berger, J. O. (1985). *Statistical Decision Theory and Bayesian Analysis*. Springer, New York, 2nd edition.

Beverton, R. and Holt, S. (1957). *On the Dynamics of Exploited Fish Populations*, volume 19 of *Fishery Investigations Series II*. UK Ministry of Agriculture and Fisheries, London.

Boyd, S., Kim, S.-J., Vandenberghe, L., and Hassibi, A. (2007). A tutorial on geometric programming. *Optimization and Engineering*, 8(1):67–127.

Breslow, N. E. (1996). Statistics in epidemiology: the case-control study. *Journal of the American Statistical Association*, 91(433):14–28.

Carey, G. and Williamson, J. (1991). Linkage analysis of quantitative traits: increased power by using selected samples. *American Journal of Human Genetics*, 49(4):786.

Chaloner, K. and Verdinelli, I. (1995). Bayesian experimental design: a review. *Statistical Science*, 10(3):273–304.

Chambers, R. and Clark, R. (2012). *An introduction to model-based survey sampling with applications*. Oxford University Press, Oxford.

Cochran, W. G. (1953). *Sampling techniques*. John Wiley & Sons, New York.

Cook, A. R., Gibson, G. J., and Gilligan, C. A. (2008). Optimal observation times in experimental epidemic processes. *Biometrics*, 64(3):860–868.

Cook, R. D. and Nachtrheim, C. J. (1980). A comparison of algorithms for constructing exact d-optimal designs. *Technometrics*, 22(3):315–324.

Darvasi, A. and Soller, M. (1992). Selective genotyping for determination of linkage between a marker locus and a quantitative trait locus. *Theoretical and Applied Genetics*, 85(2-3):353–359.

Diggle, P. and Lophaven, S. (2006). Bayesian Geostatistical Design. *Scandinavian Journal of Statistics*, 33(1):53–64.

Drovandi, C. C., Holmes, C., McGree, J., Mengersen, K., Richardson, S., and Ryan, E. (2015). A principled experimental design approach to big data analysis. http://eprints.qut.edu.au/87946/.

Drovandi, C. C., McGree, J. M., and Pettitt, A. N. (2014). A sequential Monte Carlo algorithm to incorporate model uncertainty in Bayesian sequential design. *Journal of Computational and Graphical Statistics*, 23(1):3–24.

Drovandi, C. C. and Pettitt, A. N. (2013). Bayesian experimental design for models with intractable likelihoods. *Biometrics*, 69(4):937–948.

Dykstra, O. (1971). The augmentation of experimental data to maximize $|X'X|$. *Technometrics*, 13(3):682–688.

Eidsvik, J., Mukerji, T., and Bhattacharjya, D. (2015). *Value of Information in the Eartch Sciences: Integrating Spatial Modelling and Decision Analysis*. Cambridge University Press, Cambridge.

Elliott, P., Wakefield, J., Best, N., and David J. Briggs, e. (2001). *Spatial Epidemiology: Methods and Applications*. Oxford University Press.

Goos, P. and Jones, B. (2011). *Optimal Design of Experiments: A Case Study Approach*. John Wiley & Sons.

Karvanen, J. (2009). Approximate cost-efficient sequential designs for binary response models with application to switching measurements. *Computational Statistics & Data Analysis*, 53(4):1167–1176.

Karvanen, J. (2015). Study design in causal models. *Scandinavian Journal of Statistics*, 42(2):361–377.

Karvanen, J., Kulathinal, S., and Gasbarra, D. (2009). Optimal designs to select individuals for genotyping conditional on observed binary or survival outcomes and non-genetic covariates. *Computational Statistics & Data Analysis*, 53(5):1782–1793.

Karvanen, J. and Sillanpää, M. J. (2016). Prioritizing covariates in the planning of future studies in the meta-analytic framework. *Biometrical Journal*, Accepted for publication, http://arxiv.org/abs/1608.02333.

Karvanen, J., Vartiainen, J. J., Timofeev, A., and Pekola, J. (2007). Experimental designs for binary data in switching measurements on superconducting Josephson junctions. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 56(2):167–181.

Keiding, N. and Louis, T. A. (2016). Perils and potentials of self-selected entry to epidemiological studies and surveys. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 179(2):319–376.

Kiefer, J. and Wolfowitz, J. (1959). Optimum designs in regression problems. *The Annals of Mathematical Statistics*, 30(2):271–294.

Kulathinal, S., Karvanen, J., Saarela, O., Kuulasmaa, K., and for the MORGAM Project (2007). Case-cohort design in practice – experiences from the MORGAM Project. *Epidemiological Perspectives & Innovations*, 4(1):15.

Lander, E. S. and Botstein, D. (1989). Mapping mendelian factors underlying quantitative traits using rflp linkage maps. *Genetics*, 121(1):185–199.

Langholz, B. (2007). Use of cohort information in the design and analysis of case-control studies. *Scandinavian Journal of Statistics*, 34(1):120–136.

Langholz, B. and Goldstein, L. (1996). Risk set sampling in epidemiologic cohort studies. *Statistical Science*, 11(1):35–53.

Le, N. D., Sun, W., and Zidek, J. V. (1997). Bayesian multivariate spatial interpolation with data missing by design. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 59(2):501–510.

Lindley, D. V. (1956). On the measure of information provided by an experiment. *The Annals of Mathematical Statistics*, 27:986–1005.

Lindley, D. V. (1972). *Bayesian Statistics: A review*. SIAM, Philadelphia.

Lindley, D. V. (1985). *Making Decisions*. John Wiley & Sons, London, second edition.

Little, R. J. A. and Rubin, D. B. (2002). *Statistical analysis with missing data*. Wiley, Hoboken, New Jersey.

Macgregor, S., Craddock, N., and Holmans, P. A. (2006). Use of phenotypic covariates in association analysis by sequential addition of cases. *European Journal of Human Genetics*, 14(5):529–534.

Mäntyniemi, S., Kuikka, S., Rahikainen, M., Kell, L. T., and Kaitala, V. (2009). The value of information in fisheries management: North sea herring as an example. *ICES Journal of Marine Science*, 66(10):2278–2283.

McDonald, A. and Smith, A. (1997). A tutorial on evaluating expected returns from research for fishery management. *Natural Resource Modeling*, 10:185–216.

McNamee, R. (2002). Optimal designs of two-stage studies for estimation of sensitivity, specificity and positive predictive value. *Statistics in Medicine*, 21(23):3609–3625.

Mehtälä, J., Auranen, K., and Kulathinal, S. (2015a). Optimal designs for epidemiologic longitudinal studies with binary outcomes. *Statistical Methods in Medical Research*, 24(6):803–818.

Mehtälä, J., Auranen, K., and Kulathinal, S. (2015b). Optimal observation times for multistate Markov models–applications to pneumococcal colonization studies. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 64(3):451–468.

Miettinen, K. (1999). *Nonlinear Multiobjective Optimization*. Kluwer Academic Publishers, Boston.

Miettinen, K. (2014). Survey of methods to visualize alternatives in multiple criteria decision making problems. *OR spectrum*, 36(1):3–37.

Müller, P. (1999). Simulation based optimal design. In Berger, J., Bernardo, J., Dawid, A., and Smith, A., editors, *Bayesian Statistics 6: Proceedings of the Sixth Valencia International Meeting*, volume 6, pages 459–474, USA. Oxford University Press.

Müller, P., Sansó, B., and De Iorio, M. (2004). Optimal Bayesian design by inhomogeneous Markov chain simulation. *Journal of the American Statistical Association*, 99(467):788–798.

Müller, W. G. (2007). *Collecting spatial data: optimum design of experiments for random fields*. Springer, Berlin Heidelberg, third edition.

Nagelkerke, N. J., Chunge, R. N., and Kinoti, S. N. (1990). Estimation of parasitic infection dynamics when detectability is imperfect. *Statistics in Medicine*, 9(10):1211–1219.

Nikolakopoulou, A., Mavridis, D., and Salanti, G. (2016). Planning future studies based on the precision of network meta-analysis results. *Statistics in Medicine*, 35(7):978–1000.

Prentice, R. L. (1986). A case-cohort design for epidemiologic cohort studies and disease prevention trials. *Biometrika*, 73(1):1–11.

Pukelsheim, F. (1993). *Optimal Design of Experiments*. Wiley, New York.

Raiffa, H. and Schlaifer, R. (1961). *Applied Statistical Decision Theory*. Harvard University, Boston.

Rao, J. N. and Molina, I. (2015). *Small area estimation*. John Wiley & Sons, Hoboken, New Jersey.

Reilly, M. (1996). Optimal sampling strategies for two-stage studies. *American Journal of Epidemiology*, 143(1):92–100.

Reinikainen, J. and Karvanen, J. (2016). Bayesian subcohort selection for longitudinal covariate measurements in follow-up studies. *Submitted, http://arxiv.org/abs/1609.01547*.

Reinikainen, J., Karvanen, J., and Tolonen, H. (2014). Optimal selection of individuals for repeated covariate measurements in follow-up studies. *Statistical Methods in Medical Research*, DOI:10.1177/0962280214523952.

Rezagholi, M. and Mathiassen, S. E. (2010). Cost-efficient design of occupational exposure assessment strategie–a review. *Annals of Occupational Hygiene*, 54(8):858–868.

Ricker, W. E. (1954). Stock and recruitment. *Journal of the Fisheries Board of Canada*, 11(5):559–623.

Rubin, D. B. (1976). Inference and missing data. *Biometrika*, 63(3):581–592.

Ryan, E. G., Drovandi, C. C., McGree, J. M., and Pettitt, A. N. (2015a). A review of modern computational algorithms for Bayesian optimal design. *International Statistical Review*.

Ryan, E. G., Drovandi, C. C., and Pettitt, A. N. (2015b). Fully Bayesian experimental design for pharmacokinetic studies. *Entropy*, 17(3):1063–1089.

Ryan, E. G., Drovandi, C. C., Thompson, M. H., and Pettitt, A. N. (2014). Towards Bayesian experimental design for nonlinear models that require a large number of sampling times. *Computational Statistics & Data Analysis*, 70:45–60.

Sansó, B. and Müller, P. (1999). Redesigning a Network of Rainfall Stations. In Gatsonis, C., Kass, R. E., Carlin, B., Carriquiry, A., Gelman, A., Verdinelli, I., and West, M., editors, *Case Studies in Bayesian Statistics*, number 4, pages 383–393.

Shannon, C. E. (1948). A mathematical theory of the communication. *Bell System Technical Journal*, 27:379–423, 623–656.

Susan, J. and Jack, C. (2006). Comparison of methods for analysis of selective genotyping survival data. *Genetics Selection Evolution*, 38:637–655.

Sutton, A. J., Cooper, N. J., Jones, D. R., Lambert, P. C., Thompson, J. R., and Abrams, K. R. (2007). Evidence-based sample size calculations based upon updated meta-analysis. *Statistics in Medicine*, 26(12):2479–2500.

Tokola, K., Larocque, D., Nevalainen, J., and Oja, H. (2011). Power, sample size and sampling costs for clustered data. *Statistics & Probability Letters*, 81(7):852–860.

Tokola, K., Lundell, A., Nevalainen, J., and Oja, H. (2014). Design and cost optimization for hierarchical data. *Statistica Neerlandica*, 68(2):130–148.

Van Gestel, S., Houwing-Duistermaat, J. J., Adolfsson, R., van Duijn, C. M., and Van Broeckhoven, C. (2000). Power of selective genotyping in genetic association analyses of quantitative traits. *Behavior Genetics*, 30(2):141–146.

Vanhatalo, J., Hobday, A. J., Little, L. R., and Spillman, C. M. (2016). Downscaling and extrapolating dynamic seasonal marine forecasts for coastal ocean users. *Ocean Modelling*, 100:20–30.

Vanhatalo, J., Pietiläinen, V., and Vehtari, A. (2010). Approximate inference for disease mapping with sparse Gaussian processes. *Statistics in Medicine*, 29(15):1580–1607.

Varis, O., Kettunen, J., and Sirviö, H. (1990). Bayesian influence diagram approach to complex environmental management including observational design. *Computational Statistics and Data Analysis*, 9:77–91.

Wacholder, S. (1996). The case-control study as data missing by design: estimating risk differences. *Epidemiology*, 7(2):144–150.

Wruck, L. M., Yiannoutsos, C. T., and Hughes, M. D. (2006). A sequential design to estimate sensitivity and specificity of a diagnostic or screening test. *Statistics in Medicine*, 25(20):3458.

Yokota, F. and Thompson, K. (2004a). Value of information analysis in environmental health risk management decisions: past, present, and future. *Risk Analysis*, 24:635–650.

Yokota, F. and Thompson, K. (2004b). Value of information literature analysis: a review of applications in health risk management. *Medical Decision Making*, 24:287–298.